

This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

1. (Currently amended) A system for dynamically detecting potential race conditions in a program having a plurality of threads and one or more shared memory locations, the system comprising:

with respect to each shared memory location, (i) a mechanism for maintaining a set of concurrent thread segments that access [[the]]said shared memory location, and (ii) a mechanism for maintaining a first set of locks associated with [[the]]said shared memory location;

with respect to each thread, (i) a mechanism for maintaining a set of thread segments that are ordered before [[the]]a current thread segment of [[the]]a thread, [[and]](ii) a mechanism for maintaining a second set of locks that are acquired and released by the thread, and maintaining a virtual clock associated with each thread, wherein the ordering of the set of thread segments rely on the virtual clock; and

a mechanism for reporting a warning when a potential race condition is detected.

2. (Currently amended) The system of claim 1 wherein the mechanism for maintaining the set of concurrent thread segments comprises a mechanism for maintaining a set of ordered pairs, wherein one member of a pair in the set of ordered pairs is a thread identifier, and the other member of the pair is [[a]]the virtual clock value associated with the thread identified by the thread identifier.

3. (Currently amended) The system of claim 1 wherein the mechanism for maintaining the set of thread segments that are ordered before the current thread segment comprises a mechanism for maintaining a set of ordered pairs, wherein one member of a pair in the set of ordered pairs is a thread identifier, and the other member of the pair is [[a]]the virtual clock value associated with the thread identified by the thread identifier.

4. (Currently amended) A computer-implemented method for dynamically detecting a potential race condition in a program having a plurality of threads and one or more shared memory locations, the method comprising:

with respect to each shared memory location,

maintaining a first set of locks associated with [[the]]said each shared memory location, and

maintaining a set of concurrent thread segments that access [[the]]said each shared memory location;

with respect to each thread,

maintaining a virtual clock associated with each thread, wherein the ordering of the set of concurrent thread segments rely on the virtual clock;

maintaining a second set of locks that are acquired and released by [[the]]a thread, and

maintaining a set of thread segments that are ordered before [[the]]a current thread segment of the thread.

5. Cancelled

6. (Currently amended) The method of claim [[5]]4 wherein maintaining the virtual clock comprises initializing the virtual clock to an initial value when the thread is created.

7. (Original) The method of claim 6 wherein maintaining the virtual clock comprises initializing the virtual clock to zero when the thread is created.

8. (Currently amended) The method of claim [[5]]4 wherein maintaining the set of thread segments that are ordered before the current thread segment of the thread comprises maintaining a set of ordered pairs, wherein one member of a pair is a thread identifier, and the other member of the pair is a virtual clock value.

9. (Original) The method of claim 8, further comprising, if a first thread forks a second thread:

computing the set of thread segments that are ordered before the current thread segment of the second thread as the union of (a) the set of thread segments that are ordered before the current thread segment of the first thread and (b) a singleton set containing the current thread segment of the first thread;

incrementing the virtual clock associated with the first thread, and initializing the virtual clock associated with the second thread.

10. (Original) The method of claim 9 wherein incrementing the virtual clock associated with the first thread comprises incrementing the virtual clock associated with the first thread by one.

11. (Original) The method of claim 9 wherein initializing the virtual clock associated with the second thread comprises initializing the virtual clock associated with the second thread to zero.

12. (Currently amended) The method of claim 8, further comprising:
if a first thread joins a forked thread, computing ~~the~~ a set of thread segments that are ordered before the current thread segment of the first thread as the union of:
(a) the set of thread segments that are ordered before the current thread segment of the first thread, (b) the set containing the thread segments that are ordered before the current thread segment of the forked thread but which do not belong to the first thread, and (c) ~~the~~ a singleton set containing the current thread segment of the forked thread.

13. (Original) The method of claim 12 wherein the set containing the thread segments that are ordered before the current thread segment of the forked thread but which do not belong to the first thread comprises a set containing the ordered pairs in the set of thread segments that are ordered before the current thread segment of the forked thread, such that the thread identifiers in the ordered pairs do not represent the first thread.

14. (Original) The method of claim 8, further comprising, if a thread accesses a shared memory location:

updating the set of concurrent thread segments that access the location by forming a set comprising the union of (a) a set containing the current thread segment of the thread, and (b) a set containing the thread segments in the set of concurrent thread segments that continue to access the location; and

if the updated set of concurrent thread segments contains at most one element, then updating the set of locks associated with the location to the set of locks associated with the thread, and otherwise:

- (i) updating the set of locks associated with the location to a set comprising the intersection of (a) the set of locks associated with the location and (b) the set of locks associated with the thread, and
- (ii) if the set of locks associated with the location is empty, reporting a warning of a potential race condition.

15. (Original) The method of claim 14 wherein the set containing the thread segments in the set of concurrent thread segments that continue to access the location is formed by computing a subset of the set of concurrent thread segments, wherein the subset contains each thread segment a that satisfies the following predicate:

for every thread segment b in the set of thread segments ordered before a, at least one of the following is true: (i) the thread identifier of a is not equal to the thread identifier of b and (ii) the virtual clock value of a is greater than the virtual clock value of b.

16. (Currently amended) A dynamic race detection system including a processor, wherein the system is configured to comprising:

modify and implement a compiler of a runtime system that inserts calls to a race detector in compiled code; and

implement a memory allocator of the runtime system that adds to shared memory objects instrumentation information required by the race detector.

17. (Original) The system of claim 16 wherein the compiler is a modification of another compiler.

18. (Original) The system of claim 16 wherein the memory allocator is an alteration of another memory allocator.

19. (Currently amended) A computer-implemented method for dynamic race detection, the method comprising:

by way of modifying a compiler of a runtime system, inserting calls to a race detector in compiled code; and

by way of a memory allocator of the runtime system, adding instrumentation information required by the race detector to shared memory objects.

20. Cancelled

21. (Original) The method of claim 20 wherein adding the instrumentation information required by the race detector is by way of changing the memory allocator of the runtime system.

22. (New) A system for dynamically detecting potential race conditions in a program having a plurality of threads and one or more shared memory locations, the system comprising:

with respect to each shared memory location, (i) a mechanism for maintaining a set of concurrent thread segments that access said shared memory location, and (ii) a mechanism for maintaining a first set of locks associated with said shared memory location;

with respect to each thread, (i) a mechanism for maintaining a set of thread segments that are ordered before a current thread segment of a thread, (ii) a mechanism for maintaining a second set of locks that are acquired and released by the thread, and maintaining a virtual clock associated with each thread; and

a mechanism for modifying a compiler to insert calls to a race detector in compiled code, wherein the race detector reports a warning when a potential race condition is detected.